

Disability and Adaptive Preference¹

Elizabeth Barnes

Objectivist theories of welfare are those which maintain that a person's wellbeing goes beyond the satisfaction of their desires (including both actual desires and suitably idealized versions of those desires) and their pleasure or happiness (broadly construed). According to the welfare objectivist, there is a certain thing or certain things – a state of being, a set of freedoms, functions, or capabilities, a list of basic goods, etc – that constitute the good life.² And the good life is constituted as such regardless of whether particular persons desire the objective criteria defined, or could reasonably be said to be 'happy' without them.³

An obvious *prima facie* problem for any such theory of welfare arises from the fact that many people⁴ will express preferences for lives that do not include one or more of the objective criteria laid out for personal wellbeing. And they will sometimes steadfastly persist in preferring such lives, and insisting that such lives are as or more valuable to them than ones which would meet the objective standards of wellbeing. Why shouldn't such preferences (assuming they are genuine) undermine the accuracy of whatever wellbeing criteria have been given by the objectivist?

I'm not, in this paper, concerned with the truth or falsity of welfare objectivism, or with whether the above *prima facie* problem is in fact a substantial one. Rather, I want to focus on the way objectivists *respond* to the problem. Particularly, I want to look at strategies for 'explaining away' the first person testimony of those who express preferences for lives which don't conform to a basic set of objective criteria for wellbeing. The formulation of this strategy that I'll look at in detail is that of adaptive preferences. I'm choosing the adaptive preference story because it's perhaps the most detailed and best developed account of its kind in the literature, but the concerns raised for it generalize beyond that particular case. The basic question is simply this: for two persons *x* and *y*, how can *x* get warrant for the claim that *x* knows more about *y*'s wellbeing than *y* does?

The use of the adaptive preference model to explain away certain cases of first-person testimony arises primarily from the capabilities approach to welfare, popularized by Amartya Sen and Martha Nussbaum. The capabilities approach conceives of human welfare in terms of certain basic and essential goods which are construed as *capabilities*.⁵

¹ Many thanks to Ross Cameron, Daniel Elstein, Elizabeth Harman, Gerald Lang, Daniel Nolan, Jason Turner, Georgia Testa, Christina Van Dyke, Robbie Williams, and audiences at the University of Colorado and the University of Leeds.

² Which is why they are sometimes called 'objective list theories' (see, e.g., Parfit 1984).

³ See Heathwood (2009) for an excellent overview of the classificatory issues here.

⁴ And sometimes even entire cultures

⁵ See, inter alia, Nussbaum (2001), (2001)b, (1993) and Sen (1985), (1993). The distinctive focus on *capabilities* is that theories of welfare look at what a person can do, not at what she has. Having a certain amount of resources may be vital to ensuring basic capabilities, but those resources are important only insofar as they facilitate basic capabilities. Likewise, we should aim, not toward equality of resources, but

In order to thrive, a person must be able to participate in these fundamental goods (i.e., have the *capability* to access them). If any are lacking, then that person cannot truly flourish. Likewise, in order to be considered just, a society must ensure (as far as possible) that all of its citizens have access to these basic goods.

Not all people or cultures value all the capabilities that Sen and Nussbaum have in mind, of course. But the capabilities approach has a ready and well-developed response to this problem: the psychological model of *adaptive preferences*. The adaptive preference model purports to give a compelling explanation of why a person might alter their preferences toward something which is sub-optimal. Once this explanation is in place, the thought goes, there is then no reason to take a person's preference for a life without one or more of the basic goods as *evidence* that a life without one or more of the basic goods can in fact be just as optimal as a life with them all.

In what follows, I examine the use of the adaptive preference model within the capabilities approach, which I take as a paradigm example of explaining away first-person testimony about wellbeing. As a particularly apt illustration, I'll use the case of physical disability.⁶ Physical disability represents, according to the capabilities approach, an absence of one or more basic goods (bodily integrity, physical health, etc – depending on the particular specification), and thus a life with a physical disability cannot meet the criteria of a fully flourishing human life. Yet many disabled people claim to have benefited from their experience of disability, to the extent that on the whole they prefer a disabled life to a non-disabled one. I will argue that application of the adaptive preference model to these claims is unwarranted. The reasons why such application is unwarranted will further show that the adaptive preference model is, in general, not nearly as powerful a tool as the capabilities approach often assumes it is. And though the discussion is specific to the case of adaptive preferences, the difficulties highlighted are general ones: that it is much harder to justifiably ignore or explain away first-person testimony about wellbeing than literature on the subject makes it seem.

I will begin by explaining the basics of the adaptive preference model, distinguishing the rational-choice version discussed by, *inter alia*, Elster and Bovens, from the normative version made famous by Sen and Nussbaum (section I). I will then show how the model can be applied to the case of disability (section II). After explaining the application to disability, I raise some crucial problems for it (section III). This leads to a discussion of

rather toward equality of capability. The capabilities are understood as together comprising the basic threshold for human flourishing.

⁶ A note about terminology: I will from here on use the term 'disability' to refer to a physical feature which impairs (either in preventing or greatly increasing the difficulty of) an individual's ability to carry out tasks of daily life that we generally expect persons to be able to do with relative ease. I will use the term 'disabled people' to refer to persons who have such a physical feature. Such terminology is controversial – indeed, any choice of terminology in this area will be controversial. But it's also easily translatable into other ways of speaking: for example, if you favor a social construct theory of disability, what I mean by 'disability' you mean by 'impairment'. Also note that I am explicitly concerned only with *physical* – not cognitive – disabilities. The same sorts of arguments might be applicable to the case of cognitive disability, but I haven't done the work to show that they are. In any case, we shouldn't simply *assume* from the fact that we can generically refer to both physical and cognitive disabilities with our term 'disability' that arguments which apply to the former apply to the latter, and vice versa.

whether and how one can obtain warrant for using the adaptive preference model to explain away first-person testimony about wellbeing (section IV). I end by summing up the implications of the previous sections, both for the specific case of disability, and for the capabilities approach and objectivist theories of welfare in general (section V).

I. Adaptive preference: the basics

Discussion of adaptive preference stems from the work of Jon Elster.⁷ Elster describes adaptive preferences as the phenomenon familiar from the La Fontaine fable of the fox and the grapes. The fox desires to eat a bunch of grapes hanging from a tree, but eventually realizes that they are too high up for him to reach. Having discovered that he cannot get the grapes, he decides that he didn't really want them anyway – grapes are too sour for foxes.

What has happened in this situation? According to Elster, the fox originally encounters a scenario where he believes his set of options includes getting and eating the grapes. In this scenario, he prefers to eat the grapes. But then things change: the fox finds that his set of viable options has contracted, and eating the grapes is no longer a part of it. Faced with this situation of contracted options, the fox no longer prefers to eat the grapes. For Elster, this change in preference is fundamentally irrational because it is not autonomous – it is *mere adaptation*. The fox does not *decide* to alter his preferences to suit his new situation. He simply experiences a preference change (likely from an unconscious need to avoid disappointment and frustration). The hallmark of rationality, for Elster, is autonomy – conscious decisions made in light of considered options – so preference changes that are adaptive in this respect he deems irrational.

This does not mean, however, that *all* preference changes in response to diminished options must be irrational. Elster distinguishes changes in preference which are adaptive from those which he calls 'character-planning'. Imagine an alternative fable that goes like this. The fox loves grapes, but is traveling to a country where no grapes grow, but where another sort of summer fruit is abundant. The fox decides to cultivate a taste for this new fruit, so that he will still have something tasty to eat. This new fruit is much sweeter than grapes, but after several tries and a lot of concentration, the fox realizes that this increased sweetness is in fact very pleasant – he now really does prefer the new fruit over grapes, and so could say without any hesitation that grapes are a little too sour for foxes. Elster argues that this sort of preference change – though it has the same end result and also arises from a contraction of viable options – is rational because it is the result of a conscious decision. The fox does not simply realize that he now finds grapes too sour; rather, he has made a concerted effort to cultivate a taste for a new fruit, and as a result he prefers the new fruit to grapes.

The difference between the two cases, according to Elster, is that the first change in preference arises merely from an 'affective drive' – a basic, subconscious need to not face frustration and unmet desire. The second change, in contrast, arises from a meta-desire – a desire to alter one's desires in light of new circumstances. The latter can result

⁷ See especially Elster (1983)

in autonomous, considered actions, whereas the former cannot (it is mere adaptation). So the latter is rational whereas the former is not.

Many philosophers, however, find Elster's emphasis on autonomous deliberation to be somewhat misplaced. To use an example given by Nussbaum⁸, imagine that as a child you desire to be a famous opera singer, though you later learn that you aren't much of a singer and thus being a famous opera singer isn't really within your feasible career options. It seems that it's perfectly rational for your preferences to change in light of this – for you now to no longer prefer to be an opera singer, and perhaps think that you would prefer to be, say, a philosopher (which happens to be what you are) rather than an opera singer no matter what your singing ability – even if you did not, when you were eight, consciously set about altering your preferences through the application of a meta-desire.

In response to such concerns, Luc Bovens (1992) gives an alternative way of distinguishing sour grapes-type preference changes from perfectly acceptable preference changes. Both kinds of preference change Bovens envisages can be considered 'adaptive', but one is rational while the other is not. An agent's adaptation of preferences in light of a contracted options set can be deemed rational just in case it coheres with the rest of the agent's beliefs and desires.⁹

Consider the second fox fable. This fox doesn't merely come to believe that grapes are too sour for foxes. He also comes to believe that things which are much sweeter than grapes are delicious, that anything which is the same level of sweetness as grapes is too sour for foxes, etc. That is, in adapting his preference toward the new fruit, he's not just changed his preferences about grapes, he's changed all his related taste preferences such that they cohere with his new-found preference for the new fruit.

Contrast this to the original fable. It seems no part of the fable that the fox has altered his other preferences about tastiness when he decides that grapes are too sour for foxes. He likely doesn't suddenly crave things which are much sweeter than grapes, and he wouldn't turn his nose up at an alternative tasty treat that came his way, even if it was just as sour as the grapes he claims not to want.

For this reason, Bovens argues, the fox in the classic fable looks to be doing something objectionable and irrational, whereas the fox in the alternative version of the fable does not. One fox has simply altered his preferences about grapes, which ends up in tension with many other of his beliefs, desires, and preferences. The other, in response to lack of grapes, has altered his entire preference set about tastiness. But notice that it's no part of the fox's ending up, after adapting his preferences, with a fully coherent set of taste preferences that he consciously *decided* to change his preferences. He could end up in the very same situation from having undergone a subconscious preference change, and would still be counted as having fully rational preferences. Again, Nussbaum's opera-dreamer child likely didn't one day willfully decide to change her preferences in order to

⁸ Nussbaum (2001)

⁹ For a more detailed discussion of the differences between Elster's and Bovens' views see Bruckner (2009). Bruckner argues for the rationality of an even wider range of adaptive preferences than does Bovens, though both are in agreement that Elster's constraints on rationality are far too narrow.

fit her options. Nevertheless, assuming that as an adult she has a set of preferences which are fully coherent with her preferring to be, say, a scientist rather than an opera singer, those preferences can be seen as fully rational.

Nussbaum and Sen, however, construe the issue of whether changes in preference are rational as a normative one.¹⁰ A change in one's preferences, in response to a contraction of viable options, is objectionable only in those cases where the change leaves one preferring something which is somehow less good or less optimal. That is, the change in preferences is only problematic insofar as it leads to a preference for something which one *should not*, *ceteris paribus*, prefer. A change in preference is *adaptive*, in Nussbaum and Sen's sense, just in case it represents a change in (or formation of) preferences in response to diminished options toward something sub-optimal.

Here are two paradigm examples, the first a case of adaptive preference *change* and the second a case of adaptive preference *formation*. Firstly, in the phenomenon known as Stockholm syndrome, victims of kidnapping or hostage-taking come to prefer being kidnapped – they come to believe that their kidnapper is really on a noble mission, and has rescued them, and that their kidnapping is thus of great benefit to them, etc. It's fairly easy to see how beliefs and preferences such as this could arise. The kidnap victim is put in a traumatic situation from which they see no possibility of escape, so simply in order to cope they (subconsciously or otherwise) lose the desire to escape. Such coping mechanisms may well be an admirable facet of human psychology, but we'd be very reluctant to say that the preferences of a person with Stockholm syndrome are rational, or serve as evidence that being a kidnap victim is a good way to live. Rather, according to Nussbaum and Sen, we should simply say that such preferences are adaptive.

Secondly, a woman who grows up in a deeply patriarchal society may well form preferences for submissive gender roles – shunning education, taking orders, and in some cases even accepting abuse. It's likely, though, that in her upbringing she did not consider alternatives to these – if she was even able to consider them at all – as viable options. Thus her set of viable options is limited from the start. As a result, she may well adapt her preferences to suit her situation. Again, say Nussbaum and Sen, this might be an admirable coping mechanism, but we'd hesitate to say that her preferences for submissive gender roles are rational¹¹, formed as they were by the deeply patriarchal system she was placed in, and from which she sees no escape. Thus the presence of her preference for submissive gender roles does not by itself give us any evidence that a life which coheres to submissive gender roles is just as good as a life which does not. Because her preferences are adaptive – formed toward something sub-optimal in light of a severely diminished set of options – they cannot serve as such evidence.¹²

¹⁰ See especially the discussion in Nussbaum (2001).

¹¹ Though for arguments that at least some such preferences are, see Bruckner (2009).

¹² In the way that, say, my preference for *Macbeth* compared to your preference for *Hamlet*, if we are both educated students of Shakespearean tragedy, might give some evidence (defeasible, of course) that the two are different but equally valuable plays.

In what follows, I will use the term ‘adaptive preference’ in the normatively-laden sense intended by Nussbaum and Sen, as it is this notion of adaptive preferences which is meant to help the capabilities approach.

II. Adaptive preferences: the case of disability

Now that I’ve given a basic outline of what adaptive preferences are meant to be, I’ll turn my attention to their use in defending the capabilities approach to welfare. People often express preferences for lives lacking one or more of the fundamental goods deemed necessary, by the capabilities approach, to human thriving. But, the thought goes, we should not take these claims as evidence that lives lacking those goods really can flourish if the preferences reported are *adaptive*. And since most cases of such preferences can be readily seen to be adaptive, they are not an evidential threat to the capabilities approach.

To examine how this works – and whether it is justified – I want to look specifically at the test case of physical disability. A disabled life, according to the capabilities approach, *cannot* be as optimal as a relevantly similar non-disabled life.¹³ Yet many disabled people claim the opposite.¹⁴ Are such claims easily diagnosed as adaptive preferences?

At first glance, the adaptive preference model of explanation looks particularly suited to first-person claims of some disabled people about their own quality of life. Many disabled people say that their lives are in fact enriched by their disability – that having a disability is of benefit to them and they would prefer to be disabled rather than non-disabled. Furthermore, they say that these preferences (of disabled people for a life with disability) serve as *evidence* that a disabled life is no less optimal than a non-disabled one.

These claims strike many people as odd, sometimes even incomprehensible (how could you *prefer* to be disabled?) but the capabilities approach has, via adaptive preferences, a ready explanation of them. If physical health and lack of physical impairment are – as most capability theorists maintain – on the list of features essential for human flourishing, then it is impossible to live one’s best life without them. One can be happy, certainly, but one can never truly *thrive*. A disabled life is one lacking in certain essential goods, and hence one that will always be in certain respects sub-optimal. Those who claim that

¹³ E.g., Sen includes health, mobility, and bodily integrity in ‘very elementary’ capability set, and notes that disabled people will likely have ‘less capability’ than the able-bodied (Sen 1990, 1993). Nussbaum includes in her list of capabilities ‘bodily health’, ‘bodily integrity’, and ‘senses, imagination and thought’ (see, inter alia, the appendix to Nussbaum (2001)). Her elaboration of ‘bodily integrity’ includes ‘being able to move freely from place to place’ and her elaboration of ‘senses, imagination, and thought’ includes ‘being able to use the senses’. Most disabilities – whether chronic illnesses, paralyses, non-standard bodily formations, or the absence of sense modalities – will obviously fall foul of at least one of these.

¹⁴ See, e.g., Sarah Triano’s definition of ‘Disability Pride’ from *The Encyclopedia of Disability*: ‘Disability Pride represents a rejection of the notion that our physical. . .differences from the non-disabled standard are wrong or bad in any way. . .It is a public expression of our belief that our disabilities are a natural part of human diversity, a celebration of our heritage, and a validation of our experience.’

being disabled is a good thing are thus mistaken: they misunderstand what it takes to lead a ‘life that is worthy of the dignity of a human being’.¹⁵

The diagnosis of adaptive preference behavior can explain why such mistaken judgments about one’s own life and wellbeing arise. Disabled people are placed in (or born into) circumstances which are often devastating, and from which they cannot escape. For those who acquire a disability, some physical capacity or aspect of their physical health is taken away from them, and it won’t be coming back. This often forces them to radically alter their daily habits, their life plans, etc. For those whose disability is congenital, they grow up with the knowledge that they are ‘different’ and that their life must be non-standard – that people will treat them differently, and that options most people take for granted are closed to them.

One way of dealing with such a situation, then, is to simply adapt one’s preferences. If the disabled person can manage to convince herself that that being disabled is in fact of benefit to her – that hers is a good life which it would have been reasonable to choose even if it was not forced on her – then she will be much more likely to cope well with her disability and be happy in her current situation. Such adaptive preference behavior is a useful coping mechanism, and in many cases even an admirable one. But it should not give us any evidence that her disability is *in fact* of benefit to her. Sometimes, in very non-ideal circumstances, people are forced to adapt their preferences toward what is non-ideal in order to get by. That fact alone should not give us any evidence that what is preferred in such circumstances is in fact valuable or part of a good life.

III. Problems with the adaptive preference story

1. Begging the question

The adaptive preference model, as deployed above, is explanatorily powerful. It is questionable, however, in certain key aspects. Recall the dialectic: (i) some people claim that disability is not a sub-optimal feature (i.e., it’s a way of being different, but not a way of being different such that a life with that difference is worse than a relevantly similar life without it); (ii) to support the claim of (i), they use the first-person testimony of disabled people who claim to like being disabled. The capabilities approach, in contrast, argues that (i) is obviously misguided – being disabled is not just another way of being different (the way that, say, being gay or African-American is) because having a disability entails the absence of one of the basic capabilities. The first-person reports invoked in (ii) should thus be explained away via adaptive preferences. We should not, in short, take disabled people at their word when they claim that being disabled is a good thing¹⁶, and the adaptive preference model shows us why.

¹⁵ Nussbaum (2001), pg. 5.

¹⁶ That is, when they say a disabled life is just as good as a non-disabled one. The capabilities approach doesn’t need to deny that being disabled could bring some benefit (greater fortitude, maybe) or that sometimes a person could reasonably claim that disability has *caused* an overall benefit in their life (e.g., the lonely shut-in who only develops close personal relationships once she becomes disabled – she may rightly value her new disabled life, with those relationships, more than her previous non-disabled but lonely life). But they are committed to the idea that disability doesn’t ever result in a life which is *just as good as*

But this response to the combination of (i) and (ii) is far too quick. Adaptive preferences, crucially, are modifications to one's desires to meet a set of circumstances *which are sub-optimal*.¹⁷ But whether being disabled is in fact sub-optimal is precisely the question that is up for debate. The opponent of the capabilities approach currently being considered *does not think that being disabled is something which is sub-optimal*. So to simply diagnose the first-person reports of disabled people who claim to prefer a disabled life over a (relevantly similar) non-disabled one as adaptive preferences blatantly begs the question.

To sum up: we should only diagnose adaptive preference behavior in cases where someone has modified their preferences to suit a sub-optimal situation. But whether being disabled is something sub-optimal is the question up for debate. So simply positing adaptive preferences in the case of disability begs the question.

2. Over-generalizing

The basic adaptive preference model is guilty of more than just potential question-begging. It also seems to over-generalize in striking ways. Suppose that you and I are having a debate about homosexuality: I say being gay is just a way of being different, whereas you say being gay is a sin and entails the absence of a basic human good (heterosexual relationships). I think this is crazy, and point out to you many instances of gay people thriving. But, of course, you don't perceive those examples as true human thriving, because you think in each case the person is missing a basic human good (heterosexual relationships). So I take a different approach: I point out to you several cases of first-person testimony in which gay people say that they are happy, that they enjoy being gay, and that they would not prefer to be heterosexual. Surely this is at least some evidence that these gay people are in fact having a positive experience of homosexuality? But you smile knowingly at this. *Of course* these gay people say that they are happy, that they live rich rewarding lives with loving relationships just as valuable as heterosexual ones. These reports, though, are a simple case of adaptive preference behavior. These people are put in a terrible situation – being gay! – from which they can see no escape, and so in order to cope they modify their preferences. That's an admirable psychological adaptation, certainly, but it's no evidence that a gay life is as valuable as a straight one.

Diagnosis of adaptive preference behavior, in the case above, seems clearly misguided – it's just a defense of anti-gay prejudice. But what, if anything, is the difference between the case where we apply the adaptive preference story to gay people, and the one where we apply it to disabled people?

a relevantly similar (i.e., keeping other goods fixed) life without that disability. That is, if person *a* has all the goods on the list except those ruled out by being disabled, and person *b* has all the goods simpliciter, person *b* is better off than person *a*. And if we could hold all *a*'s other goods fixed but in addition give her the goods she lacks by being disabled, we would thereby make her better off. Thanks to Daniel Nolan for discussion on this point.

¹⁷ Again, we're concerned here with the normatively-laden sense of adaptive preference employed by Nussbaum and Sen, among others.

IV. Warrant for the adaptive preference model

You might think this question should be settled independently of any discussion of adaptive preferences. A capabilities approach to welfare – any objectivist theory of welfare, for that matter – must first settle, on independent grounds, what goods should comprise the basic components of a flourishing human life. Once that list is settled, we can then proceed to discuss adaptive preferences. But we can *only* do so once the list is settled, since adaptive preferences can only be discussed in a context where we know which desires count as aiming toward the sub-optimal (those things which will detract from a flourishing life), and which do not.

Such an approach would certainly make the application of the adaptive preference model easier, but I don't think it's realistic. We don't – or at least we shouldn't – decide what constitutes a flourishing human life from the armchair. We rely heavily on experience – our own and others. We make observations and, crucially, *we listen to testimony*. Without doing so, we cannot decide what constitutes human flourishing. Yet the adaptive preference model would caution us that some of these observations are misleading, and that not all such testimony should be believed.¹⁸

To proceed at all, then, the approach needs to be holistic. We can't settle the list of basic human goods without listening to testimony about wellbeing, but if the adaptive preference model (or any other 'explaining away' strategy) is correct we shouldn't listen to all such testimony. What to do?

My discussion will proceed in terms of *warrant*. I'll argue that we need an explanation of when (and why) we are warranted in diagnosing adaptive preference behavior.¹⁹ Sometimes the adaptive preference model seems intuitively quite suitable (e.g., Stockholm syndrome), in other cases it seems hopelessly misapplied (e.g., gayness). We need a principled way of distinguishing between these two kinds of cases, to prevent the adaptive preference model from over-generalizing.

My major claim will be that, in establishing warrant for diagnosing adaptive preference behavior, the bar should be set high. Indeed, the bar should be set *very* high. The reasons for this are simple and quite general, but before discussing them it's important to get clear on exactly what's going on with the adaptive preference story. So I'll first discuss the specifics of diagnosing a case of adaptive preferences before discussing whether and how one can get warrant for such a diagnosis.²⁰

¹⁸ Case in point: Nussbaum appears to justify the list of basic capabilities by construing it a list which 'people from a wide variety of cultures. . . would choose' (2001), pg. 85. Given the list in question, however, she clearly does not conceive *disability culture* as a legitimate culture.

¹⁹ And this can be reiterated more generally: we need an explanation of when (and why) we are warranted in explaining away first person testimony about wellbeing.

²⁰ The defender of the capabilities approach might find some of the principles I list below objectionable. I think they're intuitively very plausible – and none of them are very strong – so I think it's a major cost for the capabilities theorist to deny any of them. But regardless of that, it's worth noting that what I say about warrant for diagnosing adaptive preference doesn't depend on the specifics of how I set up the dialectic. I'm setting it up this way because I think it makes the core issues clear, but the basic points that follow are

Let's start with:

(EVIDENCE) Ceteris paribus, for any person, x, x is a good source of evidence about x's own wellbeing.

This is a fairly weak principle, and one which I think most people would readily agree to. It's not saying that people are infallible sources of evidence about their own wellbeing, nor implying that the evidence they give is not defeasible. And the claim is ceteris paribus – they're not a good source of evidence if they're crazy, a pathological liar, etc. Now to be more specific:

(T-EVIDENCE) Ceteris paribus, for any person, x, x's first-person testimony is a good source of evidence about x's own wellbeing.

This doesn't follow from (EVIDENCE), but I think it's reasonable to hold it if you think (EVIDENCE) is true, and if you think testimony can ever be a good source of evidence. Much, if not most, of the evidence we get from others about their own wellbeing comes in the form of testimony. So – again, if they are genuine and sane – I think we should count that testimony as good (though defeasible, of course) evidence about their wellbeing.

But just as we naturally think that you are a good source of evidence about *your own wellbeing*, we also tend to think that:

(ARBITRARY SOURCE) Ceteris paribus, for any two arbitrary people x and y, there's no reason to think that y is a good source of evidence about x's wellbeing.

This principle seems obvious, for the simple reason that y may well know nothing about x. But it can be strengthened as follows, without being undermined:

(ARBITRARY SOURCE)*: Ceteris paribus, for any two arbitrary people x and y such that y is given basic information about x's circumstances, there's no reason to think that y is a good source of evidence about x's wellbeing.

The motivation behind the strengthened formulation of (ARBITRARY SOURCE) is simply that we know nothing about y's capacity to evaluate x, or anyone in x's circumstances. There may be prejudices, cross-cultural problems, or simple lack of understanding that make it impossible for y to reliably evaluate x's wellbeing. Thus we have no reason to think that y will be able to act as a good source of evidence. Notice, nothing in the principle implies that we have reason to think that y will be a *bad* or *misleading* source of evidence about x's wellbeing. The claim is weaker: we simply have no positive evidence to think that y will be a *good* source of evidence.

But from this, and agreement to (T-EVIDENCE), we can conclude:

independent of the set-up.

(AUTHORITY): Ceteris paribus, for any two individuals x and y, x's testimony is preferable to y's testimony as a source of evidence about x's wellbeing.

Now this is a mild claim: (AUTHORITY) does not say that x's testimony is *better* than y's. It simply says that, other things being equal, we should prefer x's testimony to y's as a source of evidence. And this seems to follow quite straightforwardly if we agree to (T-EVIDENCE) and (ARBITRARY SOURCE)*. We have reason to think that (ceteris paribus) x will be a reliable source of information about x's wellbeing, but no such reason to think that y will be a reliable source of information about x's wellbeing. So if we're faced with a choice between x's testimony and y's testimony about x's wellbeing, we should prefer (other things being equal) x's testimony.

The adaptive preference model says there are cases in which y's testimony is preferable to x's testimony. If we think the above principles are correct – and I think we should – then we should thus interpret it as singling out cases where other things are *not* equal. (AUTHORITY) tells us to prefer x's testimony in standard cases, so a case in which we prefer y's testimony will need to be non-standard.

There are two places to locate this: the adaptive preference model can say that sometimes other things fail to be equal with respect to (ARBITRARY SOURCE)* or that sometimes other things fail to be equal with respect to (T-EVIDENCE), or both. Saying that the situation fails to be ceteris paribus with respect to (ARBITRARY SOURCE)* *alone*²¹ is unpromising, and strikes of chauvinism. If (ARBITRARY SOURCE)* by itself is the locus of non-standardness (i.e., (T-EVIDENCE) is fine), then the explanation lies with y, not with x. That is, there must be something about y that makes her an unusually reliable source of information about x, such that we shouldn't prefer x's testimony *about x's own wellbeing* to y's. But this should strike us as strange: even if I am very clever, and have studied moral philosophy at great length, you will *in general* know more than I do about what it takes to make you thrive. To think otherwise would be to vastly overestimate the role of moral philosophy, and particularly the role of armchair theorizing within moral philosophy. It would be to conceive of the appropriate practice of moral philosophy as developing an a priori theory of the good and then going out into the world and applying that theory to all cases encountered, come what may. Whereas, I take it, we tend to more modestly (and practically) conceive of moral philosophy as encountering cases and then, from that basis, working out a theoretical framework that can accommodate them. Any such theorizing will of course have an armchair component to it – but its *foundation* is at least partly an empirical one.

The better option, it seems, is to locate the deviation in (T-EVIDENCE). Other things being equal, x's testimony is a good source of information about x's own wellbeing. But sometimes other things aren't equal: x might be insane, or a pathological liar. A defender of the capabilities approach needs to classify testimony in purported cases of adaptive preference behavior as of precisely this sort: unreliable, because for some reason the

²¹ Though thinking that other things aren't equal with respect to (T-EVIDENCE) will likely have the consequence that other things aren't equal with respect to (ARBITRARY SOURCE)* as well.

person in question is not to be believed. The adaptive preference model gives us an explanation for why the testimony is, in these cases, unreliable.²²

We should now be clear on the basic claim that someone needs to make in order to diagnose adaptive preference behavior. But the question remains: how does someone gain *warrant* to make such a claim? Unlike the most straightforward examples where we'd be warranted in disbelieving personal testimony, adaptive preference cases aren't ones of *general* skepticism. If you are crazy, or a pathological liar, I have good reason not to believe your testimony about anything. But if I think you're demonstrating adaptive preference behavior, I'm in the unique position of thinking I should believe your testimony about most things – just not your testimony about your own wellbeing.²³

So if y wants to diagnose adaptive preference behavior in x, y must say to the otherwise competent x that, for some special reason, y knows more than x does about x's own wellbeing. And y must do so without simply stipulating that x has sub-optimal preferences, if question-begging is to be avoided.

In some cases – Stockholm syndrome, battered women – the application of the adaptive preference model looks appropriate. But in others – e.g., gayness – it looks terrible. That is, in some paradigm cases it looks perfectly permissible to think that your own evidence of the badness, for the persons involved, of a situation trumps their first-person reports of its goodness. In many other cases, though, the directly analogous move is clearly impermissible. Whatever you may be predisposed to think about the person in that situation's wellbeing, the person is clearly a preferable source of evidence about their own wellbeing. What gives you warrant in the former cases, but not in the latter, to explain away first-person testimony via the adaptive preference model?²⁴

To address this question, I offer a modest claim and a more ambitious claim.²⁵ The modest claim should be fairly uncontroversial, but is also less interesting. The ambitious claim is more interesting, but also more tenuous.

²² Specifically: when placed in a situation that significantly limits their options in sub-optimal ways, people change their preferences in order to cope, so we should not interpret their reports of satisfied preferences as reports of *genuine* wellbeing.

²³ Or even more strongly in some cases: just not your testimony about your own wellbeing with respect to some specific feature.

²⁴ Again, what follows is a discussion of warrant for adaptive preference *as construed by the capabilities* theorist (Nussbaum, Sen, etc), where an explicitly normative component is included. It's uncontroversial, I would think, that most people who have become disabled and now prefer a disabled life have undergone a preference *change*. Likewise, people who were born disabled and prefer a disabled life have undoubtedly formed these preferences in response to being disabled (just as women who prefer to be female have formed this preference in light of being female, and gay people who prefer to be gay have formed these preferences in light of being gay). These alterations and formations of preferences certainly seem 'adaptive' in the sense used by Bovens (1992). Whether they are 'adaptive' in the sense used by Elster is likely an empirical question. The question here is whether we have warrant to say they are 'adaptive' in the normatively-laden sense.

²⁵ Again, these are set up specifically with reference to the adaptive preference strategy I've been discussing, but the basic ideas I'm arguing for should reiterate for most *any* attempt to explain away first-person testimony about well-being.

First, the modest claim:

(HARD WARRANT): warrant for the diagnosis of adaptive preference behavior is hard to get.

The reasons for this are simple. The basics of the adaptive preference model can easily over-generalize. And when the adaptive preference model does over-generalize, it's very difficult to realize that it has done so. Because of the way the adaptive preference story works, the model allows the perpetuation and justification of prejudice by dismissing the need for refinement (based on new evidence) in our moral thinking. These two considerations combined give the adaptive preference model the potential to be epistemically very damaging. It should, thus, only be used with caution. To ensure such caution, warrant for positing adaptive preference behavior should be considered hard to get.

To put it another way: the adaptive preference model over-generalizes easily, and there's good reason to think we'd be bad at detecting the cases of over-generalization.²⁶ So, for any particular case where you'd prima facie think you're warranted in diagnosing adaptive preferences, the problem of over-generalization is a defeater for that justification. So you'll need a defeater for that defeater if you're going to get your warrant back. Warrant will thus be difficult to come by.

To illustrate this, again consider the ways in which the adaptive preference model could be applied to defend prejudice and ignore evidence. Imagine the ardent sexist who reasons as follows. Some women certainly *claim* to enjoy being female, and to have benefited from their gender. But even if we believe that these claims are genuine that's no reason to think that being female is in fact just as valuable as being male. Women are put in a devastating situation: they see the superiority (both physical and mental) of maleness, yet they cannot escape their femaleness. It's natural, then, that in order to cope they modify their preferences, convincing themselves that being female is in fact a good thing for them, and that they would even in some cases prefer to be female than male. This psychological adaptation probably enables them to cope better with the daily disappointments of being female, but the testimony it produces is no evidence that femaleness is in any way on a par with maleness. We should recognize women's first-person testimony as what it is: a valuable psychological adaptation. We as men, however, know what it takes to live a truly flourishing human life, and know for independent reasons that someone without male gender can never attain such flourishing.

The argument above, of course, strikes of parody. But one need only look to history (the sorts of things said about women during the suffrage movement, for example) to realize that it's not at all far-fetched.

When the adaptive preference model goes wrong, it goes very wrong. The basic point of (HARD WARRANT), then, is one of epistemic humility and caution: you are in general a much better source of evidence about your own wellbeing than I am. So if I am going to

²⁶ The vast literature on implicit bias seems to support this. Just because we think we're operating from the light of pure reason alone doesn't mean we are.

instead claim that I am a better source of evidence than you are, I had better have very good reason to do so, considering the potential for such claims to over-generalize in ways that will simply defend prejudice and the moral status quo. Warrant for the adaptive preference model is difficult to obtain.

Now the more robust positive proposal:

(SOCIAL WARRANT): warrant for adaptive preferences can only be gotten in cases which involve a general *social* distortion.²⁷

By social distortion I mean things like: abuse of power relationships, exertion of dominance, forcible removal by one party of another party's resources or freedoms, etc. Notice first that this criterion for diagnosing adaptive preference behavior gets the paradigm cases right. Abused women and kidnap victims have both been subjected to such social distortions. Someone had an advantage over them, and used that advantage to keep them submissive and subdued. When they report having *benefited* from their experience, they are reporting to have benefited from this sort of social distortion. Contrast women and gays. Being female or being gay is not any sort of social distortion between two parties. It is nothing to do with anyone but the particular person in question – it's how they are in and of themselves. When someone reports to have benefited from being female or from being gay, they are reporting to have benefited from some fact about themselves. You can't be an abused spouse or a kidnap victim unless someone acts against you. But you can certainly be female or gay without someone acting against you – you can be female or gay on a desert island.

Another, perhaps more controversial way to capture the same distinction is to characterize it as a difference between features which are intrinsic and those which are extrinsic.²⁸ The analogous claim – call it (Extrinsic Warrant) – would then be that warrant for adaptive preferences can only be obtained for cases of a preference shift or formation toward an extrinsic feature. The idea here is that we cannot justifiably diagnose adaptive preference behavior for someone's claims about how they are in and of themselves. But you might be able to justifiably diagnose adaptive preference behavior for someone's claims about how they are extrinsically (in relation to their environment, other people, etc).

In addition to getting the paradigm cases right, it might prove easier to provide independent argument for the *badness* of certain social distortions – e.g., abuse of power relationships, etc – than for the badness of non-social or intrinsic features of a person. I think this is supported by the fact that in trying to argue for the badness of an apparently intrinsic feature, people will sometimes try to claim that it is in fact a disguised social

²⁷ NB: the claim is an *only if* not an *if and only if*. Since my project is a critical one, I'm interested only in exploring what would *have* to be in place to diagnose adaptive preference. So the claim is only that social distortion is a necessary condition – not a sufficient condition – for an adaptive preference diagnosis.

²⁸ It is more controversial because some people worry that features as rich as gender, sexuality, etc can never be fully intrinsic. But the basic thought is just this: you could be gay, female, or disabled (contra the social model of disability) if you'd grown up on a desert island. But you couldn't be a kidnap victim or an abused spouse if you'd grown up on a desert island.

distortion. The religious right, for example, uses as a cornerstone in its argument that homosexuality is a sin the claim that homosexuality is a psychological deviation that occurs in response to childhood trauma, abuse, or neglect. They then reason that the adult preferences of gay people, since they were formed in response to a deeply scarring childhood event, should not be taken as evidence for the value gayness.

Here's my claim: the religious right is *correct* about the *conditional* claim. If being gay really were just a response to painful childhood trauma, then we might have reason to question whether the preferences of adult gay people toward being gay are in fact good evidence for the value of gayness. If we in fact discovered this empirical evidence about the origin of gayness, then it would be plausible to diagnose the preferences of adult gay people as adaptive. And that this conditional is plausible gives strong support to a principle like (SOCIAL WARRANT). *If* gayness were caused by some sort of social distortion (e.g., childhood abuse), *then* we might have warrant to diagnose adaptive preference. The religious right is, of course, blatantly and reprehensibly wrong in endorsing the antecedent of the conditional. But the conditional itself is helpful to think about, when considering what it would take to get warrant for the diagnosis of adaptive preference.²⁹

To sum up: (HARD WARRANT) says we should make warrant for diagnosing adaptive preference behavior difficult to obtain; (SOCIAL WARRANT) outlines one potential way of doing this – saying that we are only warranted in diagnosing adaptive preference behavior in situations that represent some form of social distortion. I am most interested in defending (HARD WARRANT), but I think some version of (SOCIAL WARRANT) may well be a plausible specification of (HARD WARRANT).

Now let us return to the test case of disability. Some disabled people claim to enjoy their experience of disability, and even to prefer that they be disabled rather than non-disabled. A standard capabilities-based approach would have us treat these first-person reports as instances of adaptive preference behavior. But notice that whether we endorse (SOCIAL WARRANT) or the weaker (HARD WARRANT), the first-person reports of disabled people are *not* plausibly subject to an adaptive preference diagnosis. If we favor the stronger (SOCIAL WARRANT), this follows quite straightforwardly. That a person has a disability is a fact about herself, rather than a social distortion.³⁰ Being disabled, just like being gay or being female, is a way a person is in and of themselves. If we accept anything like

²⁹ Aren't disabilities sometimes caused by social distortions? Yes. But it is no part of *what it is* to be disabled that disability involves such social distortions (and of course most disabilities do not). In just the same way, *what it is* to be male does not involve social distortions simply because in some cases a social distortion (sexist-based sex selection by parents) causes the existence of a male child rather than a female child. To be disabled (or to have a specific sex or sexual orientation) is just to be a certain way in and of oneself. What we tend to think about Stockholm syndrome – and what the religious right claims about gayness – is that *what it is* to have that feature is to stand in a certain cause-and-effect relationship to a specific social distortion or class of social distortions. That is, the very nature of these features involves the social distortion in question. And that more robust sense of 'x is caused by social distortion' is what I'm claiming can give warrant for the diagnosis of adaptive preference.

³⁰ Though, of course, a social model of disability would disagree with this (claiming that disability is a social construct). I am not going to discuss the social model of disability for the purposes here, but I think the above considerations do give us reason to prefer an alternative. Regardless, the social model can agree that *impairment* is intrinsic.

(SOCIAL WARRANT), then, we can't attribute the reports of disabled people who claim to have benefited from their disability to adaptive preferences.

The weaker (HARD WARRANT) doesn't imply that same conclusion quite as straightforwardly, but it still supports it. Disabled people have a unique and vibrant culture, have strong advocacy groups, make valuable contributions to the arts (which they couldn't have made if not disabled), and, in the case of deaf persons, even have their own language. These are all social markers of a legitimate minority group, and each makes the dismissal of persons who claim to enjoy being disabled as merely demonstrating adaptive preferences increasingly implausible.

Disabled people still, of course, have to contend with basic skepticism: how on earth could you *prefer* to be disabled? I think that a major part of this skepticism, though, comes from the way we tend to view disability.³¹ Blindness, for example, is just the absence of sight. I can shut my eyes and get a sense of what it would be like to go blind – no fun at all. Disabled people, however, tell a different story. Rebecca Atkinson, for example, writes of her experience of going blind gradually:

Losing your sight is not like just shutting your eyes. The loss is so gradual that as one sense dies others grow. Suddenly you can smell the world and sense when someone is standing out of your line of vision. Your brain grows on the inside and things on the outside start to matter less. I get to live my life twice over in two different bodies (the sighted one I used to have and the partially sighted one I now have), and with that comes the privilege of spying on the world and its intricacies from multiple vantage points. . . . If this experiment of going blind has taught me anything, it is that what you lose in one place you gain elsewhere, and while a blind life is different to a sighted life, it is not lesser.³²

And in a similar vein Helen Keller describes her experience of blindness and deafness:

My world is built of touch-sensations, devoid of physical color and sound; but without color and sound, it breathes and throbs with life. Every object is associated in my mind with tactual qualities which, combined in countless ways, give me a sense of power, of beauty, or of incongruity: for with my hands I can feel the comic as well as the beautiful in the outward appearance of things. Remember that you, dependent on your sight, do not realize how many things are tangible.³³

The experience of disability, for many disabled people, is not just one of absence (of a sense modality, a function, physical health, etc). It is, rather, one of absence in particular areas that creates (*in virtue of that very absence*) opportunities in other areas – opportunities not open to the non-disabled. And some³⁴ disabled people report that the

³¹ See Barnes (forthcoming) for extended discussion of the subsequent issues.

³² Atkinson (2007)

³³ From 'The Seeing Hand' in Keller (2003)

³⁴ Not all of course – just as not all women are happy being women and not all gay people are happy being gay. Some are, some aren't. The most plausible characterization of any of these traits is that they are

resulting experiences disability creates mean that, on the whole, disability is of great benefit to them.

In light of this, it seems that even the weaker (HARD WARRANT) plausibly rules out applying the adaptive preference story to disability. Disabled people give reasonable (if surprising) justifications for why their lives are benefitted by disability, and they have unique contributions to culture, the arts, and even language to back them up. That's more than enough, I'd argue, to put disability outside the scope of the adaptive preference model, if we're assuming something like (HARD WARRANT).

V. Conclusions

There are several conclusions to draw from the above discussion. Firstly and most specifically: on an appropriate construal of the adaptive preference model, it does not undermine the first person testimony of disabled people. The more vocally they argue for rights and respect, the more disabled people publically make claims about disability that many non-disabled people find surprising and sometimes implausible. But if the account of warrant for adaptive preferences I've given above is accurate, then we cannot simply dismiss those claims as instances of adaptive preference behavior.

Secondly: the capabilities approach does not deal adequately with the phenomenon of physical disability. The capabilities approach claims to advocate for disabled people, because it has a way of showing why society must meet their unique needs (by way of societal obligation: societies are required to provide access to as many of the basic goods as they can, and bring all their citizens as close as possible to a life that has access to all the goods, so they must provide disabled people with the resources to recover as many of the basic goods as they can³⁵). But their method of advocating for disabled people requires that we assume that a disabled life is sub-optimal when compared to a non-disabled one, and in order to do this, that we willfully ignore the first-person reports of many disabled people. They claim that the adaptive preference model explains why this is justified, but as I have argued above this claim is suspect. The capabilities approach cannot justifiably ignore these first-person claims, and thus does not have an adequate account of disability. Indeed, without further argument it presents a view which may well be construed as *prejudiced* against disabled people.

Thirdly and more generally: even if we grant to the capabilities approach that the adaptive preference model is coherent and tenable, its application will be much more

neutral – neither such that they make you automatically worse off, nor such that they make you automatically better off. How you react to them will depend on what else they're combined with (including other traits, and outside circumstances). See Barnes (forthcoming) for discussion.

³⁵ E.g., where bodily integrity is compromised, they must provide mobility aids, wheelchair ramps, etc in order to restore as much of that integrity as possible, even if this means devoting a large amount of resources to them (see, inter alia, Sen 1990, 1995, 2004). Where health is compromised, they must provide excellent medical care. Where sense modalities are compromised, they must provide viable alternative (instruction in Braille, ASL, etc). In most cases, these measures will not get a disabled person to a place where they can participate fully in the basic goods, but they will get them *closer*. It will make their lives *better* even if it does not make their lives optimal or fully flourishing. This is part of a strategy which Anita Silvers characterizes as approaching 'disability rights as compensatory rights' (Silvers 2005).

restricted than much of the discussion of it has suggested. The case of disability is an apt one: capabilities theorists want to place freedom from physical impairment on their list of basic goods, but the adaptive preference model, if appropriately constrained, cannot assist them in doing so. Indeed, it's far from obvious that the adaptive preference model, if used with evidential restraint, can support anything like the basic list of goods which capabilities theorists want to defend. So even if we grant that the basic idea behind the adaptive preference model is legitimate, it's questionable whether the model can do the work which the capabilities approach assumes it does.

Finally, and most broadly: these sorts of concerns will generalize. Any objectivist theory of welfare needs to be able to provide a story about why a person's actual wellbeing sometimes conflicts with their perceived wellbeing. The adaptive preference model is one such story, the one used by the capabilities approach. And thus insofar as the above discussion is specific to the adaptive preference model it is specific to the capabilities approach (and other closely related theories). But the *kind* of concerns raised above will, I think, emerge for many objectivist theories of welfare. Again, the case of disability is apt. We are happy, in moral philosophy, to discuss from the armchair how (if there is a choice) it is always morally obligatory to bring a non-disabled person rather than a disabled person into existence, whether and how it could ever be permissible to knowingly bring a disabled person into existence at all, etc.³⁶ Most of the people making such claims are aware of the disability pride movement. It's just that the claims of disabled activists are easy to dismiss – they are adaptive preference, they are wishful thinking, they are false hope, they are misguidedly person-affecting preferences, they are simply confused, etc.³⁷ The general moral of the above discussion, I think, should be that many of these attempts to explain away first-person testimony (of disabled people, in this instance, but the phenomenon emerges many places³⁸) are far too quick. That's not an argument in favor of welfare subjectivism. It is, rather, the simple claim that any defense of an objectivist theory of welfare needs to be more cautious and more rigorous in explaining away first-person testimony about wellbeing, especially the first-person testimony of persons who take themselves to be a legitimate and valuable minority group.

³⁶ See, inter alia, Parfit (1984), Kamm (2002), Harman (2004), (2007), McMahan (2005), Hare (2007)

³⁷ See, inter alia, Brock (2005), Kahane (forthcoming), McMahan (2005), (2005)b. Harman (this volume) develops one of the more detailed versions of the 'explaining away' strategy. Harman argues that the first-person preferences of the disabled can be construed as strongly person-affecting (e.g., they are happy with their disabled lives, and don't identify with the people they would have been had they not been disabled). But, she argues, the presence of such person-affecting preferences shouldn't influence what we think in the objective case (e.g., just because someone is glad she is deaf doesn't mean that we should think that a deaf life is, for her or for others, just as good as a hearing one). Though Harman's paper is explicitly concerned with specific arguments against curing deafness (and for most of her conclusions she does not need the assumption that being deaf is worse than being hearing), such concerns about person-affecting preferences could potentially be used to undermine the claim that the first-person testimony of disabled people counts as *good evidence* about their objective wellbeing. Again, though, the crucial question seems to be whether we can get warrant for thinking – without begging the question or over-generalizing – that disability is something bad or sub-optimal toward which merely person-affecting preferences have been directed. (Compare, e.g., the case in which I characterize someone's preference for a gay life as merely person-affecting, and run a similar argument.)

³⁸ And can, of course, readily be seen in history, in our moral reasoning about gays, women, racial and ethnic minorities, etc.

Without such caution, we run the risk of never being able to learn new and surprising information about what a good life can look like.

Works Cited

- Atkinson, Rebecca (2007): 'Do I Want My Sight Back?' *The Guardian*, Tuesday 17 July, 2007.
- Barnes, E (forthcoming): 'Disability, minority, and difference'. *The Journal of Applied Philosophy*.
- Brock, Dan (2005): 'Preventing Genetically Transmitted Disabilities While Respecting Persons with Disabilities' in Wasserman, et al eds. *Quality of Life and Human Difference: Genetic Testing, Health Care, and Disability*. Cambridge: Cambridge University Press, 87-100.
- Bruckner, D (2009): 'In defense of adaptive preferences'. *Philosophical Studies* (142), 307-24.
- Bovens, L. (1992): 'Sour grapes and character planning.' *The Journal of Philosophy* (89), 57-78.
- Elster, J. (1983): *Sour Grapes: Studies in the subversion of rationality*. Cambridge: Cambridge University Press.
- Hare, C. (2007): 'Voices from another world: must we respect the interests of people who do not, and will never, exist?'. *Ethics* (117), 498-523.
- Harman, E. (this volume): "'I'll Be Glad I Did It" Reasoning and the Significance of Future Desires'. *Philosophical Perspectives*.
- . (forthcoming): 'Harming as Causing Harm' in Roberts and Wasserman, eds *Harming Future Persons*. Springer.
- . (2004): 'Can We Harm and Benefit in Creating?' *Philosophical Perspectives* 18, 89-113.
- Heathwood, C (2009): 'Welfare' in J. Skorupski (ed.) *The Routledge Companion to Ethics*. London: Routledge.
- Kahane, G. (forthcoming): 'Non-identity, self-defeat, and attitudes to future people'. *Philosophical Studies* (available via on-line first).
- Kamm, FM (2002): 'Genes, Justice, and Obligations to Future People'. *Social Philosophy and Policy* 19. 360-88.
- Keller, Helen (2003): *The World I Live In*. NY: The New York Review Books Classics.

- McMahan, Jeff (2005): 'Causing Disabled People to Exist and Causing People to Be Disabled'. *Ethics* 116. 77-79.
- (2005)b: 'Preventing the existence of people with disabilities' in David Wasserman, Jerome Bickenbach, and Robert Wachbroit, eds., *Quality of Life and Human Difference: Genetic Testing, Health Care, and Disability*. Cambridge: Cambridge University Press. 142-71.
- Nussbaum, M.C. (2001): 'Symposium on Amartya Sen's philosophy: 5 adaptive preferences and women's options. *Economics and Philosophy* (17), 67-88.
- (2001)b: *Women and Human Development: The Capabilities Approach*. Cambridge: Cambridge University Press.
- (1993): 'Non-relative virtues: an Aristotelian approach' in Nussbaum and Sen eds. *The Quality of Life*. Oxford: OUP. 242-69.
- Parfit, Derek (1984): *Reasons and Persons*. Oxford: OUP. 351-79.
- Sen, A. (1985): *Commodities and Capabilities*. Oxford: OUP.
- . (1990): 'Justice: Means vs. Freedoms'. *Philosophy and Public Affairs*, 19 (2), 111-21.
- (1993): 'Capability and well-being' in Nussbaum and Sen eds. *The Quality of Life*. Oxford: OUP. 30-53.
- . (1995): *Inequality Reexamined*. Boston: Harvard University Press.
- . (2004): *Rationality and Freedom*. Boston: Harvard University Press.
- Silvers, Anita: 2005: 'People with Disabilities' in LaFollette, ed *The Oxford Handbook to Practical Ethics*. Oxford: OUP.
- Triano, Sarah (2005): 'Disability Pride' in G. Albrecht et. al. (eds) *The Encyclopedia of Disability*. Chicago: SAGE Publications.